www.bsc.es

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Reproducibility of EC-Earth

Towards a protocol for CMIP6

F. Massonnet, M. Ménégoz, M. Acosta, X. Yepes, O. Bellprat, M. Asif, E. Exarchou, K. Serraddell, O. Mula-Valls, V. Guemas, F. Doblas-Reyes

Barcelona, 26 may 2016

Rationale: how reproducible are climate simulations?

- If CMIP6 has to start today, can we distribute the EC-Earth model on platforms that have very different characteristics?
- Are we underestimating the hardware uncertainty?
- Can we exchange restart files from center to center?

All these important (and equivalent) questions can only be answer if a strict protocol is developed, and applied everytime a new model version is available.

This is the goal of this presentation.

Review of existing literature

- Baker et al., GMD, 2015
  - Designed to test portability of CESM model
  - Intermediate (1-yr) ensemble simulations + PCA to explore changes.
- Thomas et al., Wea. Forecast., 2002
  - Regional high-resolution atmospheric model
  - Very short (weather, 48h) ensemble simulations
- Hong et al., Mon. Wea. Rev., 2013
  - Global atmospheric model
  - Short (medium-range, 10 days) ensemble simulations
- Düben and Palmer, Mon. Wea. Rev. 2015
  - Global atmospheric model
  - Intermediate (~1.5 yr) ensemble simulations
- Other unpublished work:
  - Servonnat et al., note about reproducibility of IPSL model
    - Century long, one-member simulation
  - Janakiraman et al., conference paper.
    - Atmospheric model
    - Short (5 day) ensemble simulations

Defining a protocol for global climate simulations (1/3)

- 20-yr long, 5-member, pre-industrial, coupled simulations

Allows to look at impact of machine on mean state/bias (not possible in the case of 1-yr simulations)

Allows to measure differences due hardware as compared to internal variability

Working under stationary conditions removes possible dependence of hardware impact on the mean state

Addresses the problem from a global point-of-view; suitable to give recommendations for CMIP6

- EC-Earth3.1 is used. Note:
  - Compiling with **`-O2 -g -traceback -vec-report0`**: model runs
  - Compiling with **`-O2 -fp-model precise -fimf-arch-consistency=true -no-fma -g -traceback -vec-report0 -r8`** model runs
  - Compiling with **`-O2 -fp-model precise -fimf-arch-consistency=true -no-fma -g -traceback -vec-report0 -r8 -fpe0`**: model cannot run (crashes after 3 time steps)

Defining a protocol for global climate simulations (2/3)

- Exactly the same NEMO, IFS and OASIS codes are used. The same compilation options are used for the model code. Compilation options are **`-O2 -g -traceback -vec-report0`** except in the sensitivity experiments labeled “options” where they are **`-O2 -fp-model precise -fimf-arch-consistency=true -no-fma -g -traceback -vec-report0 -r8`**
- The number of processors used is the same (72) except in sensitivity experiments labeled « HighProc » where 512 are used.
- Initial conditions are the same, except in NEMO: a white noise of 10^{-4} K is added to generate the five ensemble members (same perturbation for all machines)
- The same version of Autosubmit (workflow manager) is used to ensure the exact same management of experiments across machines
- The libraries (NetCDF, GRIBEX, GRIBAPI, etc.) are the default ones on each machine – they are from different versions and have not necessarily been compiled with the same options

Defining a protocol for global climate simulations (3/3)

- The set of metrics of Reichler and Kim (BAMS, 2008) is used to identify where changing the hardware could have led to a difference in essential oceanic and atmospheric variables
- Whenever a difference is to be detected, a Kolmogorov-Smirnov test is done to detect if differences between two hardware configurations are systematically greater than internal variability.

Machine-dependence of the mean state

ECMWF = CCA machine
MN3 = MareNostrum3 machine
Ithaca = Ithaca (IC3) machine

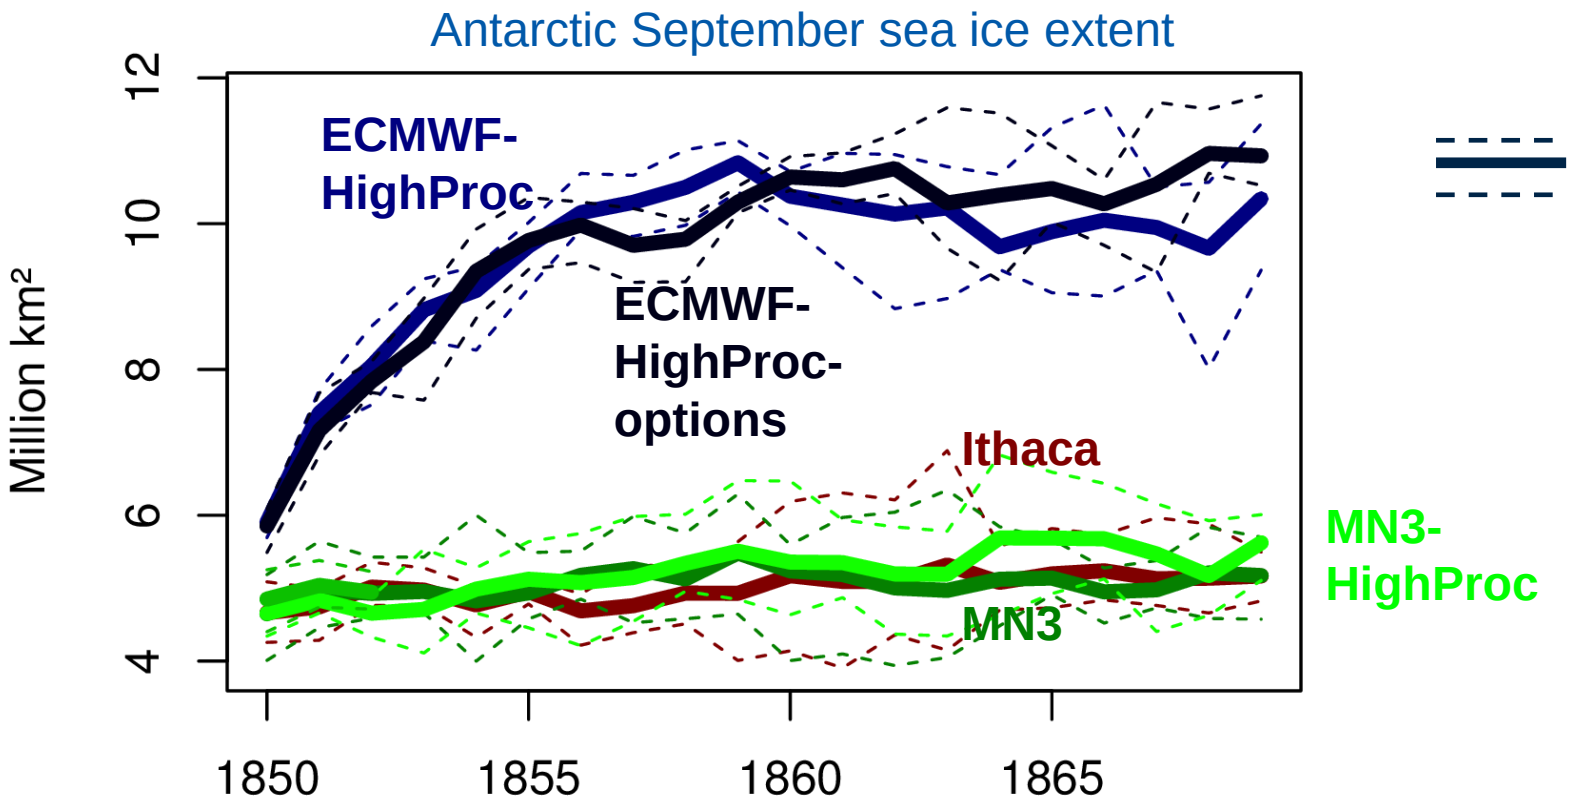

Machine-dependence of the mean state

- All other things being equal, EC-Earth3.1 is insensitive to number of processors and compilation options
- EC-Earth3.1 is sensitive to change in platform. Possible reasons:
  1) There are uninitialized arrays in the code; depending on the machine, these arrays are filled with whatever is in the memory at that time
  2) There is a bug, like division by zero, that is not caught since –fpe0 is not enabled. Values resulting from the division by zero are interpreted differently depending on the compiler
  3) Loose compilation options in the libraries (NetCDF, GRIB, …) on one platform cause reading/writing errors in the model, and this not appear in the other machine that uses more strict options.
- But remember that the –fpe0 option is *not* enabled, as this caused the model to crash on either machine

Changing the machine affects the Southern Hemisphere climate

Difference of mean T2M for MN3-HighProc minus ECMWF-HighProc. Stippling = significant at 5%

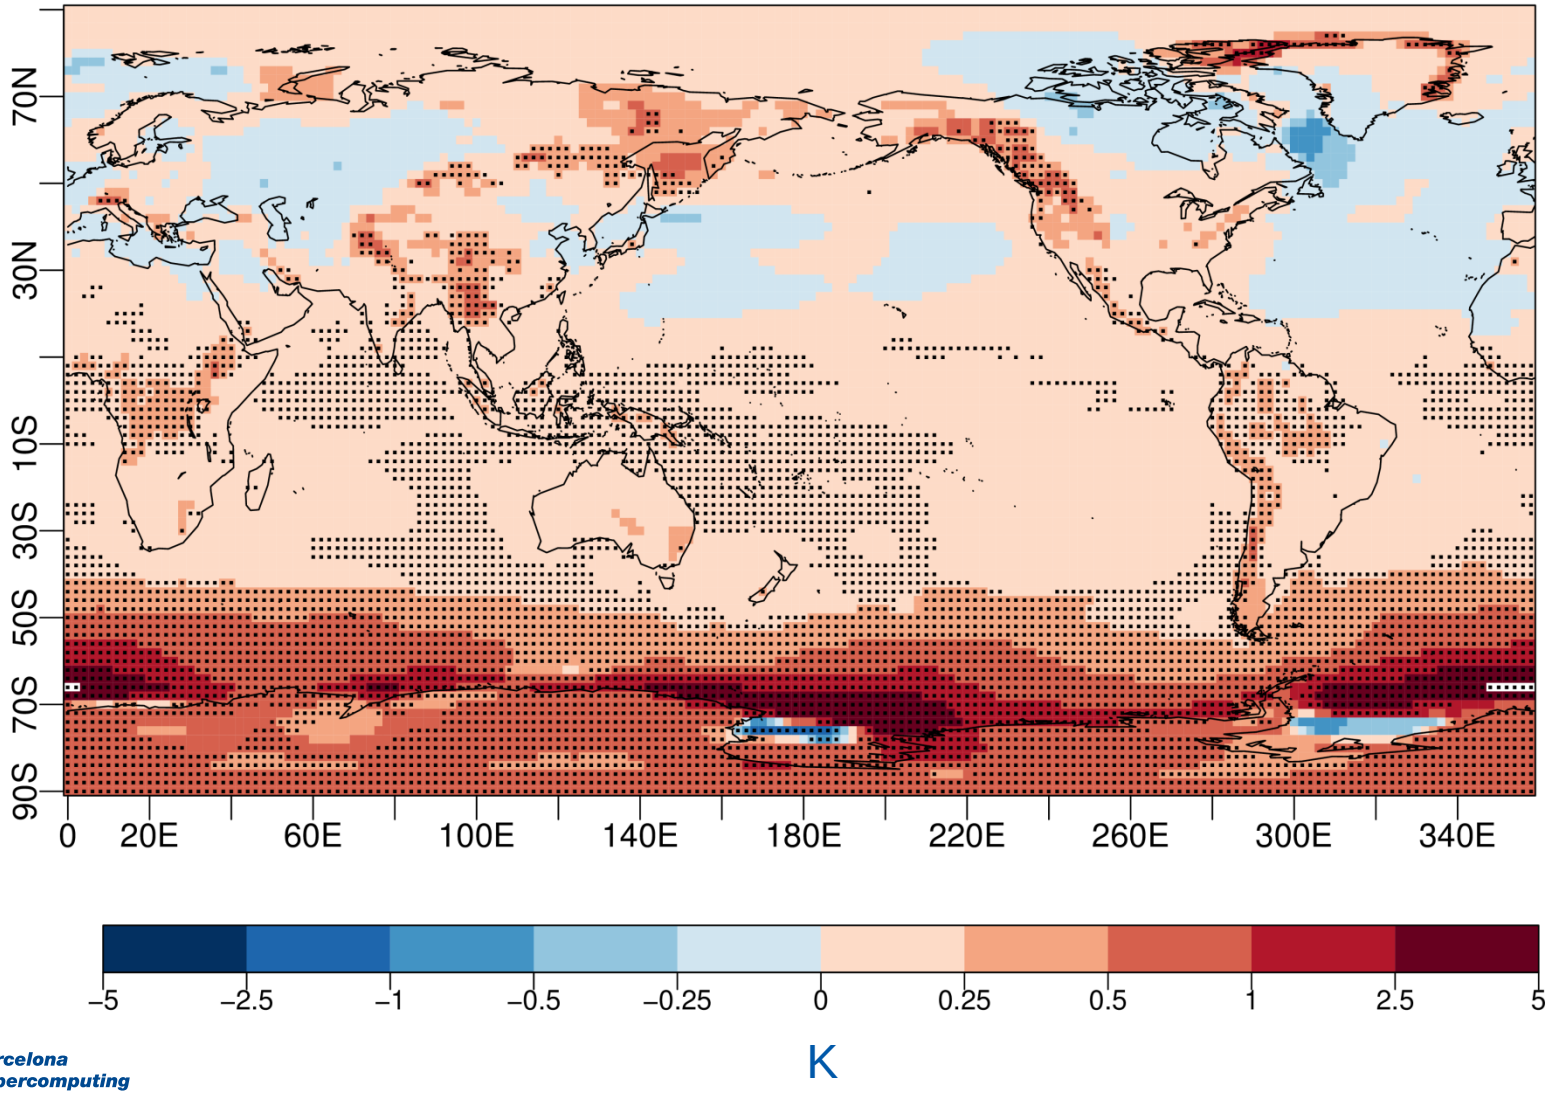

Origin of the differences: freshwater forcing from Antarctica?

Difference in SSS, after 15 days of simulation, ECMWF-HighProc minus MN3-HighProc [PSU]

Standard deviation of SSS, after 15 days of simulation, of the 5 members of ECMWF-HighProc [PSU]

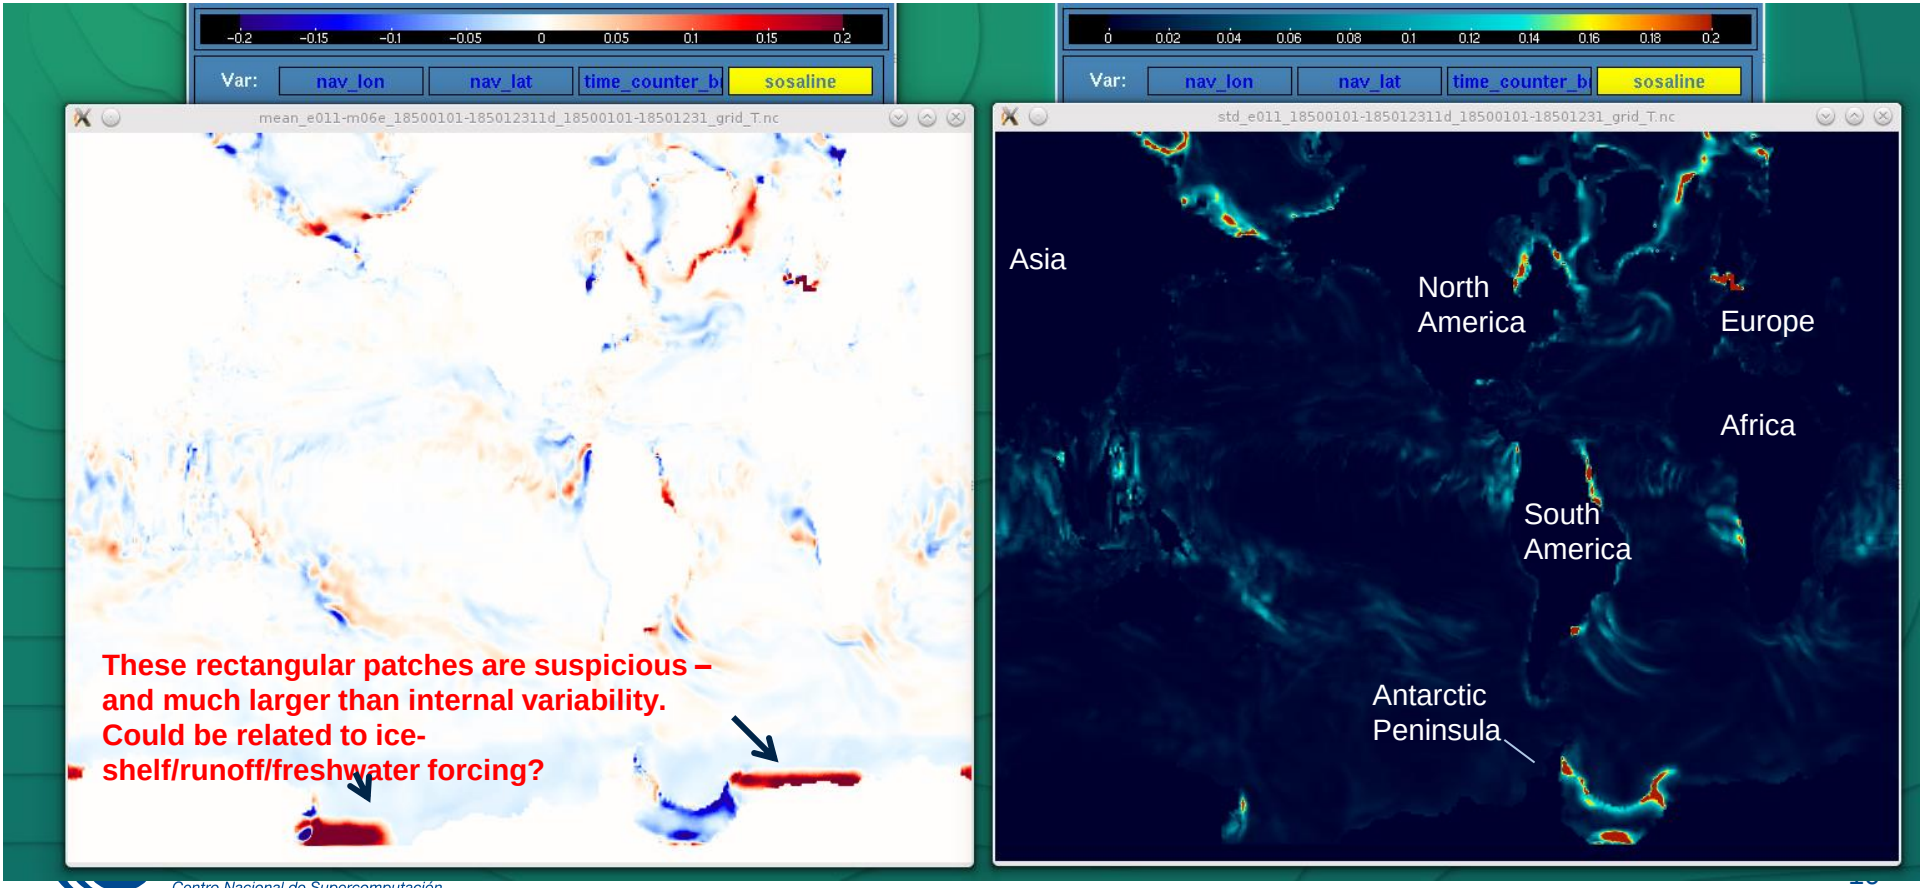

Changing the number of processors only does not affect the results

Difference of mean T2M for MN3-HighProc minus MN3. Stippling = significant at 5%

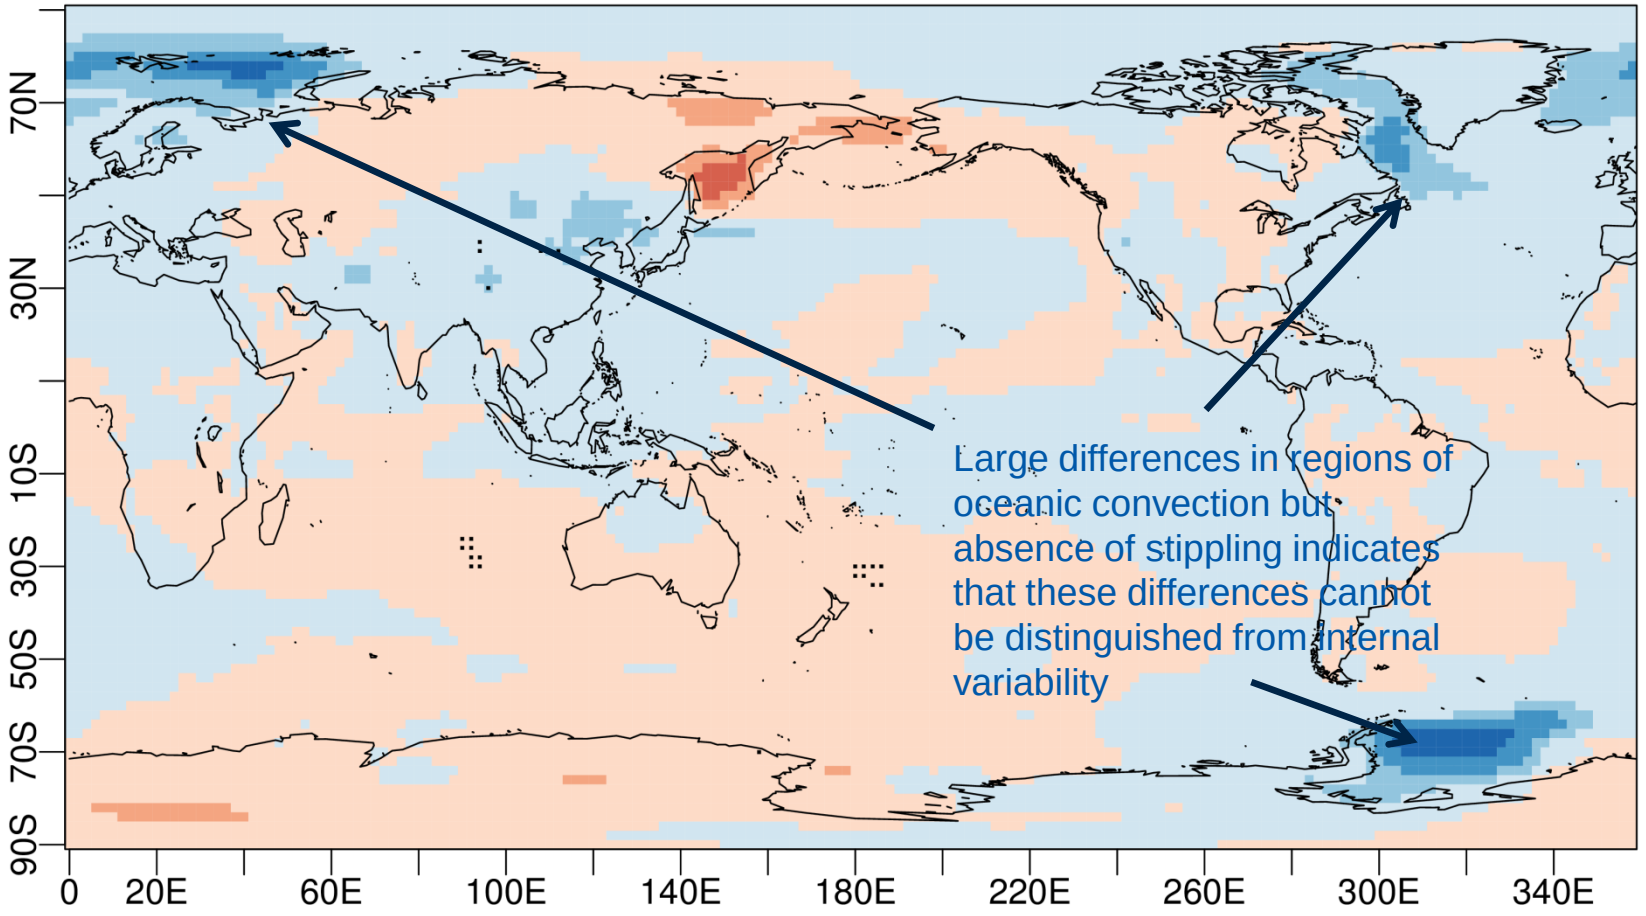

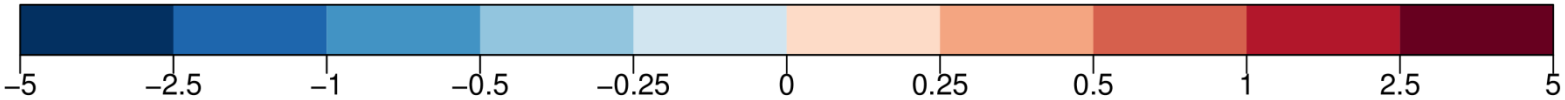

Center
Centro Nacional de Supercomputación

Changing compilation options does not affect the results

Difference of mean T2M for ECMWF-HighProc minus ECMWF-HighProc-Options. Stippling = significant at 5%

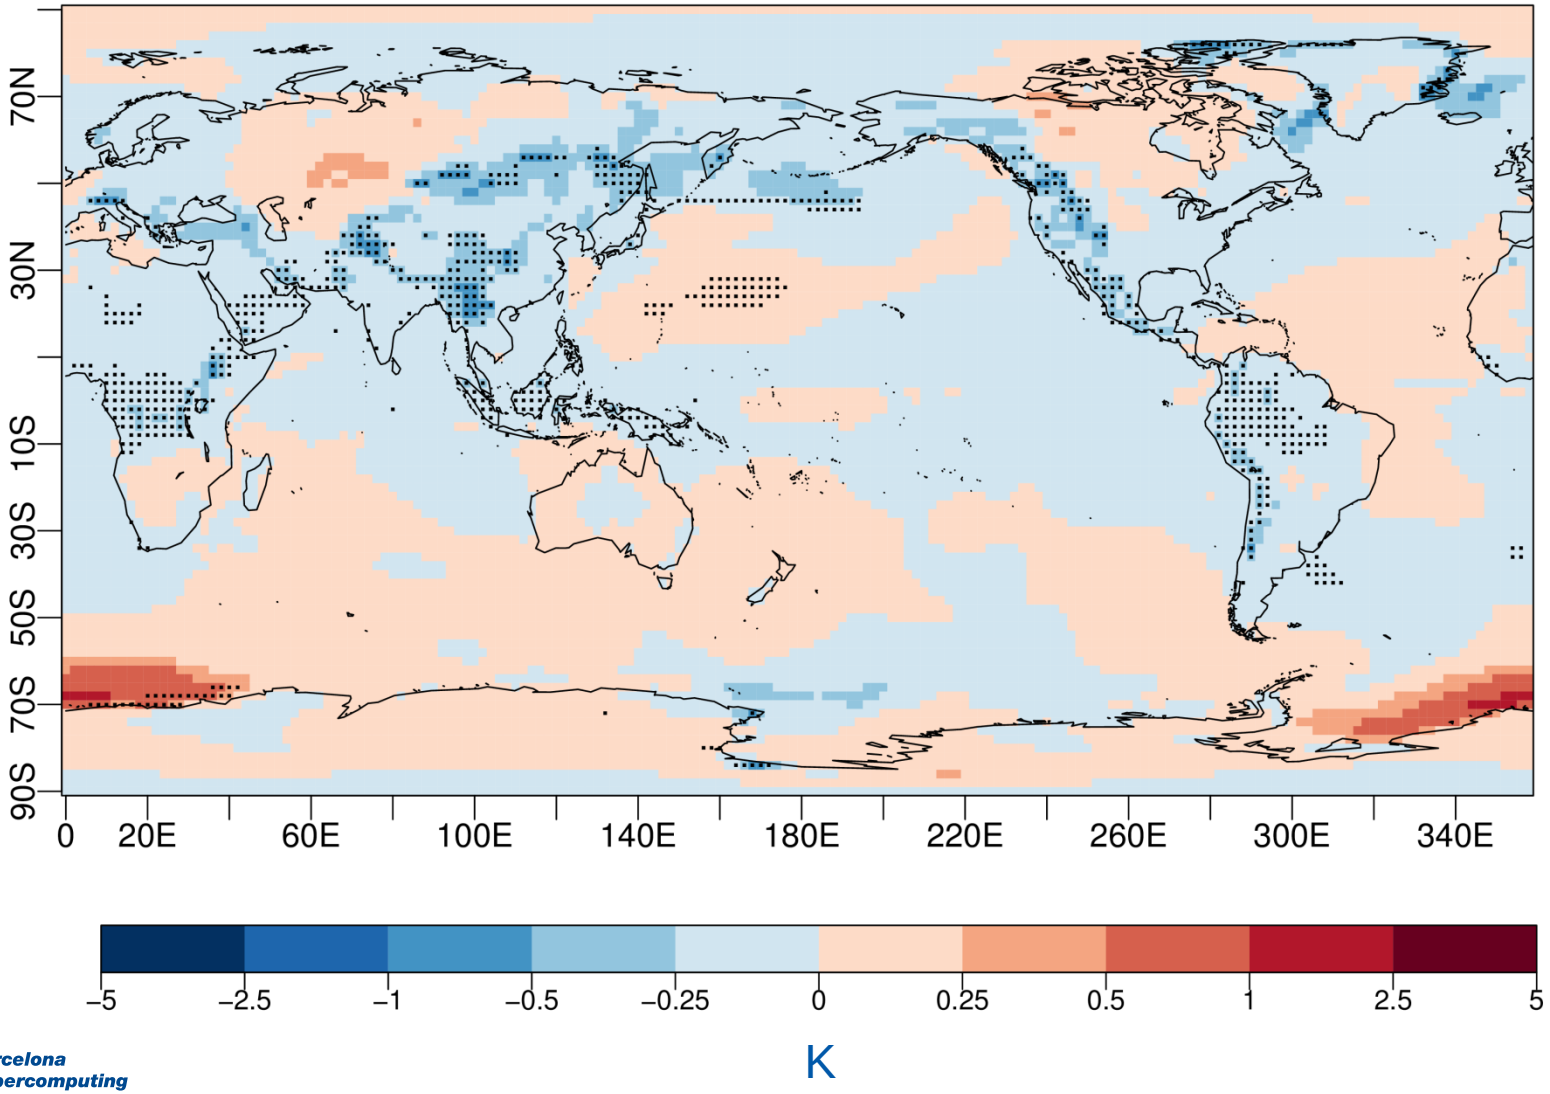

Conclusions and outlooks

- A strict protocol has allowed to objectively highlight when EC-Earth3.1 is portable and when it is not.
  - The model climate in the Southern Hemisphere climate is sensitive to the platform used
  - The model climate is insensitive to # of procs and compilation options
- The –fpe0 flag (to catch floating-point exceptions) was *not* enabled in these simulations; otherwise the model crashes.
  - New tests are now realized with EC-Earth3.2beta and –fpe0 enabled.
- The version of libraries, as well as compilation options differ likely from machine to machine.
  - Should these libraries be compiled together with the model, whenever a new experiment is started?
- **Repeating such a protocol of reproducibility is key if the EC-Earth community wants to distribute the load of CMIP6 simulations among partners – current results indicate reproducibility (from a climate sense) is far from straightforward.**